

Maximizing Privacy in MIMO Cyber-Physical Systems Using the Chapman-Robbins Bound

Rijad Alisic, Marco Molinari, Philip E. Paré, and Henrik Sandberg*

Abstract—Privacy breaches of cyber-physical systems could expose vulnerabilities to an adversary. Here, privacy leaks of step inputs to linear-time-invariant systems are mitigated through additive Gaussian noise. Fundamental lower bounds on the privacy are derived, which are based on the variance of any estimator that seeks to recreate the input. Fully private inputs are investigated and related to transmission zeros. Thereafter, a method to increase the privacy of optimal step inputs is presented and a privacy-utility trade-off bound is derived. Finally, these results are verified on data from the KTH Live-In Lab Testbed, showing good correspondence with theoretical results.

I. INTRODUCTION

Digitalization is rapidly transforming many aspects of society, using data collected by sensors in smart cities, manufacturing facilities and energy networks, in order to decrease costs, detect faults and improve the experience of its end-users. However, the introduction of these sensors makes eavesdrop attacks possible, breaching the confidentiality of the system that is being spied upon. Eavesdropping is difficult to detect, since it does not affect the system directly. The leaked information could be used by the attacker to figure out the structure of the underlying system and learn its weaknesses. A way to keep the information from getting into the hands of an adversary could be by means of encryption. However, encryption comes with its own set of difficulties, for example, increasing time delays in data streams or increased maintenance cost due to secret key handling [1]. Additionally, the increased processing time could make the system more susceptible to *Denial-of-Service* attacks, since it becomes easier to make the system miss its real-time computation constraints.

Instead, another defense strategy would be to introduce noise into the data stream, which makes the adversary uncertain about what the actual signal is. An example of where this approach is used is in the concept of differential privacy, which is a popular tool that hides user information in databases [2]. A database can release various structures of its data to anyone without explicitly revealing its individual data entries. However, an adversary could combine this information with side information to deduce the individual entries, thus breaching the privacy of the database. A differentially private database removes this possibility, for example, by corrupting the answers to queries with noise so that it is not possible to reveal individual entries with additional side information.

*Rijad Alisic, Philip E. Paré, and Henrik Sandberg are with the Division of Decision and Control Systems and Marco Molinari is with the Division of Applied Thermodynamics and Refrigeration at KTH Royal Institute of Technology, Sweden (e-mail: rijada, marcom, philipar, hsan@kth.se)

In dynamical systems the data has an additional component, namely time. The adversary can make use of models of the system in order to reconstruct corrupted data. One definition of privacy in this context is the estimation error of the system's internal states which is proposed in [3]. Introducing noise is a central component in that work as well and increasing the estimation error variance is related to increasing differential privacy [4]. A more direct generalization of differential privacy to dynamical systems is shown in [5], where the privacy of input signals is considered.

Privacy has also been considered in the context of hypothesis testing [6]. An attacker considers the value of the state of the system to be different hypotheses, and uses the measurements to determine which hypothesis is true. The privacy is defined as being the type-II error of a hypothesis test, namely the probability of missing to declare that the correct hypothesis is true.

Guided by an example of a privacy leakage scenario in a multi-residential smart building, we consider privacy to be a combination of concepts from [3] and [5]. A smart building uses sensors to read its current state, for example, temperature and CO₂ in different rooms. It also uses actuators, via a controller, to shift the building system into states which are desirable for its residents. The desirable states depend on what rooms the residents are currently in. Thus, there are streams of data inside the building which are being transmitted between the sensors, the control system, and the actuators. These data streams contain confidential information about the residents, for example, if they are home or not.

Assume now that there exists an attacker who gains access to some of the sensor-to-controller data streams, with the objective of detecting changes in the controller-to-actuator data streams, which indicates that a resident has moved from one room to another. Let the attacker know the model of the system, its initial state, and the shape of the input sequence *a priori*. However, the attacker does not know at what time step the input sequence starts and, therefore, the attacker wants to estimate it. The estimation of the change time is done using a series of hypotheses tests, which translates into solving a *change point problem*. Although the literature on this type of problem is extensive, there are no *uniform minimum variance estimators* (UMVE) for detecting when a change occurs [7], [8]. Instead, the adversary is forced to solve a combinatorial problem. Since there are no UMVE for the change time of a step input, we define privacy to be the lowest obtainable variance of the estimated change time using any estimator.

This work is an extension of a previous paper [9], where

the problem was restricted to single-input-single-output systems. There, an analysis of the system structure was conducted in order to figure out what structural components enable privacy leaks. For example, it was shown that a large signal-to-noise ratio (SNR) was a contributing factor. However, a defender can easily control the SNR by modifying the additive noise in the measurements. Instead, the analysis showed how slow dynamics also produce a large privacy leakage, since it allows for the adversary to collect more samples of the output signal during the change. The results of that study also revealed that unstable and integrator systems have the largest privacy leaks.

In this paper, we consider a similar setting, but with the extension to multiple-input-multiple-output (MIMO) systems. The inputs are assumed to be step changes, so that the system reaches some desired steady states, which implies that we only consider stable systems with no integrator dynamics appearing in the output. In this setup, fundamental lower bounds on the estimation of the change time are derived. We show that in some cases it is possible to change the inputs to satisfy the same constraints at steady state, while simultaneously increase the difficulty of estimating the change time for an adversary.

In the next section, we formulate the problem by posing two main research questions. The first question asks how much information about the change time an adversary can uncover and the second asks if there is a way to use pre-existing noise in the system to change the inputs so that privacy leaks are reduced. The answers to these problems are given in Section III, and their implications are analyzed. Finally, in Section IV, we return to the motivating example of an adversary eavesdropping on a smart building. There we highlight what the theory predicts about an adversary's capabilities of breaching the privacy of the residents and how a defender could design input signals that minimize the privacy leakage.

II. PROBLEM FORMULATION

Consider a linear time invariant system where the measurements are corrupted by a zero-mean, stationary, white Gaussian signal, e_k , with covariance $\mathbb{E}[e_k e_k^T] = \Sigma_e, \forall k$. Then, the system model M can be written as:

$$M: \begin{cases} x_{k+1} = Ax_k + Bu_k \\ y_k = Cx_k + e_k, \end{cases} \quad (1)$$

where $x_k \in \mathbb{R}^n$, $y_k \in \mathbb{R}^m$, $e_k \in \mathbb{R}^m$ and $u_k \in \mathbb{R}^p$. The system matrices $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times p}$ and $C \in \mathbb{R}^{m \times n}$ together with the noise model for e_k define M . Denote the sequence of outputs and inputs as $Y = \{y_k\}_{k=0}^N$ and $U = \{u_k\}_{k=0}^N$, respectively. The input sequence, U , is assumed to be a step,

$$u_k = \begin{cases} 0, & \text{for } k < k^*, \\ u, & \text{for } k \geq k^*, \end{cases}$$

where $\|u\|_2$ is the size of the input, and $u/\|u\|_2$ is its direction.

The objective of the attacker, is to estimate the change time, k^* , using the model M and the measurements Y . A

defender's main purpose is then to make it as difficult as possible for the attacker to obtain their goal. Motivated by the attacker's goal, we define privacy in the following manner:

Definition 1. Consider an estimator of the change time k^* for the inputs $U = \{u_k\}_{k=0}^N$, which are fed through system M in (1). Denote the estimator of k^* by ψ , which has a bias that is bounded. We define the privacy of system M to be the lowest achievable variance of the estimated change time,

$$\min_{\psi} (\text{Var}(\psi|k^*)).$$

This definition of privacy is general and the defender may consider estimators which take very complex information into account. The problem of interest in this paper, however, is to calculate the privacy of system M , conditioned on the type of estimators that the attacker can produce.

Problem 1. Let an estimator of the change time k^* in U , denoted by $\psi_u(Y, M)$, have access to the model (1) and the measurements Y of length N such that $N \geq k^*$. What is the minimum variance that any such estimator can achieve?

In physical systems, measurement noise is typically present or, alternatively, injecting noise into measurements might come with some costs. Consequently, a defense strategy might include ways to choose the input such that the existing noise is used to hide the input. Therefore, we pose the following question:

Problem 2. Consider an estimator of k^* , $\psi(Y, M)$, that has access to the model (1) and the measurements Y of length N such that $N \geq k^*$. Is it possible to choose U so that the lower bound on the variance of any such estimator is increased?

The answer to these questions will show what structures in the model, M , expose the change time, k^* , to an adversary. The answer also provides the defender with information about how to design their system so that estimating the change time becomes as difficult as possible. Although any level of privacy can be achieved by injecting enough noise into the system, additional noise also degrades the controller performance. If the controller aims to minimize a cost function, then the actual cost increases when noise is added. It is therefore important that the noise which is already present is used to the fullest extent, which could be done by placing the sensors strategically or by designing controllers that make multiple actuators cooperate so that a particular change is more difficult to estimate. In this paper, we will assume that the defender knows the noise model *a priori*, which might not always be true for real systems.

III. MAIN RESULTS

The Cramér-Rao lower bound [10] is typically used to answer questions like Problem 1. A difficulty here is that the Cramér-Rao lower bound is only defined for continuous parameters, whereas k^* takes discrete values. Therefore, a more general result is required in order to answer Problem 1. Such a result is made possible by the Chapman-Robbins (CR) bound [11]. With the CR-bound, it is possible to show what structures in the model M expose the change

time to an adversary. Therefore, subsequently we will use this information to find an input that increases the smallest possible variance of the estimate, defined as $\mathcal{B}_u(M)$.

Theorem 1. Consider any estimator of the change time k^* in the input sequence U . Denote the estimator by $\psi(Y, M)$ with bias $g(k^*)$, where M is a MIMO-system. Then

$$\text{Var}(\psi_u(Y, M)|k^*) \geq \mathcal{B}_u(M), \quad (2)$$

where,

$$\mathcal{B}_u(M) := \max_{\tau} \frac{(\tau + g(k^* + \tau) - g(k^*))^2}{e^{u^\top \mathcal{S}(\tau, M) u - 1}},$$

for $\tau \in \{1, \dots, N - k^*\}$. Here,

$$\mathcal{S}(\tau, M) = \sum_{k=k^*+1}^N \left(C \tilde{A}(k, \tau) B \right)^\top \Sigma_e^{-1} C \tilde{A}(k, \tau) B, \quad (3)$$

where,

$$\tilde{A}(k, \tau) = \begin{pmatrix} \sum_{l=k^*}^{\min(k^* + \tau - 1, k - 1)} A^{k-1-l} \end{pmatrix}. \quad (4)$$

Proof. See the Appendix. \square

Much attention will be given to $\mathcal{S}(\tau, M)$ in the subsequent sections, since it determines the smallest possible variance of any estimator, $\mathcal{B}_u(M)$. It will be assumed that the estimator is unbiased, so that $g(k^*) = 0$. The minimum variance in (2) depends explicitly on u , which implies that different inputs provide different levels of privacy.

Definition 2. Let $\lambda_{\tau}(\mathcal{S})$ be the smallest eigenvalue of $\mathcal{S}(\tau, M)$. The most private input direction is defined to be the eigenvector u_* , corresponding to the λ_{τ^*} eigenvalue, where τ^* maximizes (2).

Furthermore, it may be impossible to estimate the change time for some particular directions of u . These directions are given by the following definition.

Definition 3. We say that u_* is a fully private input direction if $\lambda_{\tau^*}(\mathcal{S}) = 0$.

Definition 2 is justified by the following proposition, which states that step-changes in the direction of u_* provide the most privacy to a system.

Proposition 1. The smallest variance for estimating change time of the step, k^* , is maximized in the direction of u_* .

Proof. Minimizing $u^\top \mathcal{S}(\tau, M) u$ for $\tau = \tau^*$ maximizes the bound in (2). \square

Proposition 1 provides a method to find the most private input directions. Therefore, it also gives an affirmative answer to the question that is posed in Problem 2.

Notice that so far in the discussion, we have not mentioned the size of the inputs. Theorem 1 states that a larger $\|u\|$, implies a smaller $\mathcal{B}_u(M)$ and thus less privacy. Therefore, it is fully possible that the change time in the the most private input direction, u_* , is easier to estimate than some u , which is not in that direction, if $\|u_*\| > \|u\|$. This is related to the SNR, which was treated in [9].

A. Fully Private Input Directions

Fully private input directions have an interesting property, namely that they make $\mathcal{B}_u = \infty$. Therefore, it is impossible to estimate when step-changes in this direction occur, since any unbiased estimator of k^* will have an infinite variance. Methods for finding fully private input directions, will be presented in this subsection.

Theorem 2. A fully private input direction exists if and only if

$$\text{rank}(\mathcal{O}) < p, \quad (5)$$

where

$$\mathcal{O} = \begin{bmatrix} CB \\ CAB \\ CA^2B \\ \vdots \\ CA^{N-1}B \end{bmatrix}. \quad (6)$$

Before we prove this theorem, we need to relate the null space of (3) to (6).

Lemma 1. The following holds for any $\tau, \tilde{\tau} \in \mathbb{N}$,

$$\mathcal{O}(\tilde{\tau})u = 0 \iff \mathcal{O}(\tau)u = 0,$$

where

$$\mathcal{O}(\tau) = \begin{bmatrix} C \tilde{A}(0, \tau) B \\ C \tilde{A}(1, \tau) B \\ C \tilde{A}(2, \tau) B \\ \vdots \\ C \tilde{A}(N-1, \tau) B \end{bmatrix}.$$

Proof. Without loss of generality, we can set $\tilde{\tau} = 1$, where $\mathcal{O}(1) = \mathcal{O}$ in (6). The proof for

$$\mathcal{O}u = 0 \implies \mathcal{O}(\tau)u = 0,$$

is obtained by using (4), which shows that $C \tilde{A}(k, \tau) B$ is a linear combination of $CA^l B$, $\forall l \in \{1, \dots, N-1\}$. The proof for

$$\mathcal{O}u = 0 \Leftarrow \mathcal{O}(\tau)u = 0,$$

follows from by rewriting (4) as,

$$\tilde{A}(k, \tau) = \begin{cases} I, & k=0, \\ A^k + \tilde{A}(k-1, \tau), & 1 \leq k \leq \tau-1 \\ A^{k+1-\tau} \tilde{A}(\tau-1, \tau), & \tau \leq k \leq N-1. \end{cases}$$

Thus each $C \tilde{A}(k, \tau) B$ is a linear combination of $CA^k B$ and $C \tilde{A}(l, \tau) B$, for all $l < k$, which implies

$$0 = \mathcal{O}(\tau)u \implies \mathcal{O}(\tau)u = \mathcal{O}u. \quad \square$$

Let us return to the proof of Theorem 2.

Proof of Theorem 2. If (5) holds, then by Lemma 1 there is a u_* such that

$$\mathcal{O}u_* = \mathcal{O}(\tau)u_* = 0,$$

which implies that $u_*^\top \mathcal{S}(\tau, M) u_* = 0$ for all τ , making $\mathcal{B}_{u_*}(M) = \infty$.

If $\mathcal{B}_{u_*}(M)=\infty$, then there is a τ^* such that

$$u_*^\top \mathcal{S}(\tau^*, M) u_* = 0 \quad (7)$$

for some u_* . Since Σ_e is positive definite, (7) implies that

$$C\tilde{A}(k, \tau^*) B u_* = 0, \forall k.$$

Finally, the last step is obtained by applying Lemma 1. \square

Theorem 2 states that there exists a fully private input direction if the input-observability matrix, \mathcal{O} , is rank deficient. If it is rank deficient, then the fully private input direction is in the null space of \mathcal{O} . Inputs in this direction do not affect the output, which is similar to what inputs that excite zero dynamics do. For sample horizons that are large enough, namely $N-k^* > n$, fully private input directions are a special case of these. Recall the definition of a transmission zero:

Definition 4. A zero, z_0 , is a complex number that makes the Rosenbrock system matrix rank deficient,

$$\text{rank} \left(\begin{bmatrix} A - I z_0 & B \\ C & 0 \end{bmatrix} \right) < m + n, \quad \text{where } m \geq p.$$

We denote x^0 as the zero-state direction and u^0 as the corresponding zero-input direction, where

$$0 = \begin{bmatrix} A - I z_0 & B \\ C & 0 \end{bmatrix} \begin{bmatrix} x^0 \\ u^0 \end{bmatrix}. \quad (8)$$

A simpler way to determine if an input direction is fully private is by checking the zero-state direction.

Corollary 1. Let the measurement horizon satisfy $N-k^* > n$. Then an input direction is fully private if and only if it is a zero-input direction, u^0 , and

$$\begin{bmatrix} C \\ CA \\ \vdots \\ CA^{n-1} \end{bmatrix} x^0 = 0. \quad (9)$$

Proof. Note that, using the Cayley-Hamilton theorem, (8) is equivalent to,

$$\begin{bmatrix} C \\ CA \\ \vdots \\ CA^{n-1} \end{bmatrix} x^0 + \begin{bmatrix} 0 \\ CB \\ \vdots \\ \sum_{k=0}^{n-2} z_0^{n-2-k} CA^k B \end{bmatrix} u^0 = 0. \quad (10)$$

If the input direction is fully private, then the second term in (10) is zero. One may then choose $x^0=0$ which gives that (10) is zero. If the input is a zero-input direction and (9) holds, then (10) being zero implies $\mathcal{O}u^0=0$. \square

Thus, if the number of samples, N , is larger than the number of states, n , one may study the zeros and the respective zero-state direction. If the zero-state direction is in the observable space of the system M , then there exists a fully private input direction and it is parallel to the zero-input.

B. Privacy-Utility Trade-off

It might not always be desirable to use the most private input direction, especially if it is fully private. A system designer might require that some amount information about the inputs leaks in order to guarantee control performance, run diagnostics or to detect attacks. Therefore, it is useful to add something in between a ‘‘fully private’’ and a ‘‘non-private’’ input. In order to do this, a type of privacy measure is required.

Definition 5. We say that the input direction u_* for model M , is more private than u if,

$$\mathcal{B}_{u_*}(M) > \mathcal{B}_u(M), \quad \text{for } \|u_*\| = \|u\|.$$

Since the defender knows how the inputs change between different modes of operation, it can tailor the additive measurement noise by using Theorem 1 in order to hide these changes. However, there could be a cost associated with generating noise, for example, if a battery with finite energy is used to perturb the measurements. If the system already has inherent measurement noise, then the designer could create controllers that use the existing noise to mask the input.

Assume that a designer of system M wishes to create a step input which minimizes the cost at steady state,

$$\begin{aligned} & \underset{x, u}{\text{minimize}} && J(x, u) = x^\top Q x + u^\top R u \\ & \text{subject to} && x = Ax + Bu, \\ & && C_1 x = r, \end{aligned} \quad (11)$$

where $C_1 \in \mathbb{R}^{q \times n}$, $q < p$. Both $Q \in \mathbb{R}^{n \times n}$ and $R \in \mathbb{R}^{p \times p}$ are positive definite. The first constraint states that the system is at steady state, whereas the second constraint represents the desired reference value of some linear combinations of the states. The number of rows in C_1 is smaller than the number of inputs in order to ensure non-trivial solutions to the program.

Assume now that the designer is willing to pay a price in the optimal cost in order to increase the privacy. They can do so by adding a regularizer to the optimization in the following way,

$$\begin{aligned} & \underset{x, u}{\text{minimize}} && J_p(x, u) = x^\top Q x + u^\top R u + \mu u^\top \mathcal{S}(\tau^*) u \\ & \text{subject to} && x = Ax + Bu, \\ & && C_1 x = r, \end{aligned} \quad (12)$$

where τ^* is the τ that maximizes (2) for the input which minimizes (11). Additionally, let (x_*, u_*) and (x_p, u_p) denote the solutions to (11) and (12), respectively. The regularization parameter μ determines how large of a cost increase is tolerated in order to improve privacy. It is important to gain some intuition on how to choose this parameter and what guarantees can be given for a specific value of the parameter.

Theorem 3. The nominal privacy gain, δ , is lower bounded by the utility loss, $\varepsilon \geq 0$ and upper bounded by the current privacy cost,

$$\mu u_*^\top \mathcal{S}(\tau^*) u_* \geq \mu \delta \geq \varepsilon, \quad (13)$$




Fig. 1: The figure shows how τ^* in (2) changes as a function of u for a system with five states and two inputs.

where,

$$\varepsilon = J(x_p, u_p) - J(x_*, u_*),$$

and

$$\delta = u_*^\top \mathcal{S}(\tau^*) u_* - u_p^\top \mathcal{S}(\tau^*) u_p.$$

Proof. The first inequality is trivial, due to the definition of δ . The second inequality is obtained by comparing the regularized costs, J_p , and rearranging the terms:

$$\begin{aligned} J_p(x_*, u_*) &\geq J_p(x_p, u_p) \\ \Leftrightarrow x_*^\top Q x_* + u_*^\top R u_* + \mu u_*^\top \mathcal{S}(\tau^*) u_* \\ &\geq x_p^\top Q x_p + u_p^\top R u_p + \mu u_p^\top \mathcal{S}(\tau^*) u_p \\ \Leftrightarrow \mu (u_*^\top \mathcal{S}(\tau^*) u_* - u_p^\top \mathcal{S}(\tau^*) u_p) \\ &\geq x_p^\top Q x_p + u_p^\top R u_p - x_*^\top Q x_* - u_*^\top R u_*. \end{aligned}$$

The last two lines is the second inequality stated in (13). \square

One may use the bounds in Theorem 3 to choose a μ which fulfills some guarantees. For example, if a maximum cost increase, $\bar{\varepsilon}$, is tolerated, then choosing the following μ ensures that the cost increase is upper bounded,

$$\mu = \frac{\bar{\varepsilon}}{u_*^\top \mathcal{S}(\tau^*) u_*} \Rightarrow \bar{\varepsilon} \geq \varepsilon.$$

This bound is not tight, which means that the designer can tune the parameter μ in order to increase privacy and thus increase the nominal ε until $\bar{\varepsilon}$ is reached, if the initial μ does not provide sufficient privacy.

The second inequality in Theorem 3 gives an interpretation of what the regularization parameter does. By rewriting the inequality in the following manner,

$$\mu \geq \frac{\varepsilon}{\delta}, \quad (14)$$

one may see that μ limits the maximum privacy-utility trade-off. When the designer chooses a specific μ , they set the maximum tolerable utility loss per privacy unit that is gained. The bound also states that increasing the utility cost will increase the privacy as well, and conversely, if there is no privacy gain, then there will be no increase in the utility cost.

Notice that the bound in (14) is only tight if τ^* minimizes (2) for both u_* and u_p . Consider the simple illustrating

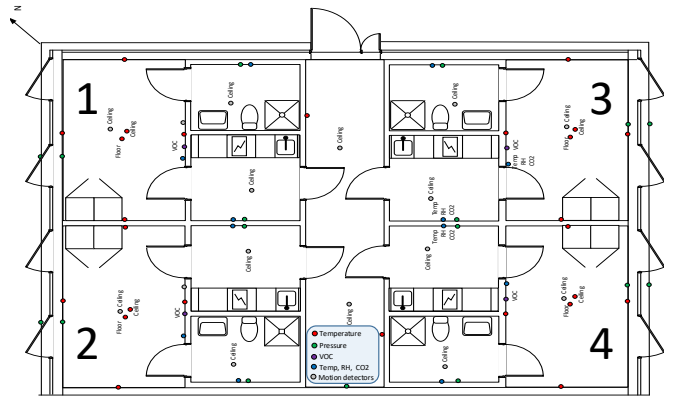


Fig. 2: An overview of the apartments in the KTH Testbed.

TABLE I: The table shows the lower bound, \mathcal{B}_u , and the empirical variance \hat{V} in different rooms. The last column shows the input projected onto the most private input direction, u_* .

Room	\mathcal{B}_u [min ²]	\hat{V} [min ²]	$u \cdot u_*$
Living Room	169	1570	1
Kitchen	5.67	277	0.002
Bathroom	18.6	145	0

example which is shown in Fig. 1, where τ_3 minimizes (2) for u_* . The regularizer in (12) will push the solution u_p towards the dashed, blue line, which gives maximum privacy for τ_3 . For some values of μ , the solution u_p might end up in a region that is maximized by τ_2 . In this case, μ can be interpreted as a looser bound on the maximum privacy-utility trade-off since the privacy gain is larger than what is captured by δ .

IV. NUMERICAL RESULTS

Let us return to the motivating example which was described in Section I, where an adversary tries to detect changes in the occupancy of different rooms in an apartment. The data for this example was taken from the IDA ICE 4.8 simulator of the Live-In Lab, KTH Testbed [12], [13]. The Live-In Lab is a physical multi-residential building outfitted with sensors in every apartment that measure temperature, relative humidity, and CO₂. The KTH Testbed is a modifiable subsection of the Live-In Lab, containing 120 square metres of living space that are split up into four apartments in this example, which can be seen in Figure 2.

The adversary is assumed to have obtained the same model of the system as the defender, for example, through studying input-output data of similar apartments. Now, let the adversary eavesdrop on the system by sampling temperature and CO₂ measurements of the different rooms in the apartment every 9 minutes. It then estimates the input change by using a Moving Horizon Estimator [14]. At the same time, the defender injects noise into the measurements in order to increase the variance of the adversary's estimation of the change time. The variance in the CO₂ sensor is four times larger than the temperature sensor, due to CO₂ being a much stronger indicator of occupancy.

TABLE II: The table shows how changing the maximum privacy-utility trade-off, μ , affects $\frac{\epsilon}{\delta}$, \mathcal{B}_u and \hat{V} .

μ	$\frac{\epsilon}{\delta}$	\mathcal{B}_u [min ²]	\hat{V}^2 [min ²]
0	-	201	6 470
$2 \cdot 10^{-9}$	$0.4 \cdot 10^{-9}$	489	45 500
$4 \cdot 10^{-9}$	$1.13 \cdot 10^{-9}$	653	732 000

A. Detecting Occupancy Changes in Apartments

Table I shows the variance of the estimated change time for different rooms in the apartment. In the last column, the projection of the input to the system onto the most private input direction, u_* , is shown. Although the empirical variance, \hat{V} , increases as the input u becomes more parallel to u_* , the same is not true for the theoretical lower bound, \mathcal{B}_u . The lower bound \mathcal{B}_u in Table I is largest for the input in the most private direction, which verifies Proposition 1, however, the input which is perpendicular to u_* does not produce the lowest theoretical bounds. This discrepancy is explained by the non-convexity of (2). As u is slowly rotated towards u_* , a different value of τ might become the minimizer of (2), with a different $\mathcal{S}(\tau, M)$. Because of this change, the input might pass a couple of local minima during the rotation. Therefore, the high empirical variance \hat{V} of the change time in the kitchen might be due to the sub-optimality of the estimator which is used.

B. Private Steady State

Let us now consider the case where the different rooms in one apartment cooperate in order to reduce the privacy leak. Consider a user entering their apartment. Then, instead of only heating the room that the user enters, the building could increase the heat production in some of the other rooms as well, thus obfuscating the attacker's estimation of the change time. This control input is obtained by solving (12). In Table II, the impact of changing the regularization parameter is shown. In the first row, the controller aims to only minimize the utility cost, whereas in the other two rows, the controller signal aims to both minimize the utility cost and to reduce the privacy leak. The private controllers are obtained for two different values of μ , under the same measurement noise covariance. Higher values of μ did not produce any noticeable improvements in the privacy. One may see that the privacy-utility trade-off bound, given by (14), holds for all instances. As discussed in Section III-B, the bound on the privacy-utility trade-off becomes tighter for smaller μ . By increasing the μ parameter, a larger trade-off is allowed. Additionally, one may see that both the theoretical and empirical variance, \mathcal{B}_u and \hat{V} , increase, as μ increases as well.

V. CONCLUSIONS

This paper shows how Gaussian noise can be used to hide the input changes to a multi-input-multi-output system. Specifically, the relation between private inputs and transmission zeros was analysed. Additionally, instead of injecting additional noise to improve privacy, a new approach where the defender makes use of the existing noise was presented. Using a convex program with a regularization term, the inputs at steady state could be made more private

by increasing the regularization parameter. The value of this parameter was shown to capture the privacy-utility trade-off. Furthermore, it was shown that increasing the steady state cost in the more private direction is a sufficient but not necessary condition for decreasing the privacy leakage. Finally, these results were verified on numerical simulations.

Connections between the Cramér-Rao lower bound and differential privacy have previously been discussed in [15]. Since the Chapman-Robbins bound is a generalization of the Cramér-Rao bound, one would expect that similar connections exist between differential privacy and (2). In fact, differential privacy can be used to establish a looser lower bound, which will be explored further in future work.

Process noise is another type of noise which affects the system, and thus is of interest for future work. The difficulty that arises in this setting is that the process noise affects multiple time steps, making the corresponding expression for (3) much more complex. Another future research direction is to analyze the same situation under smoother changes of the input. This alternative approach would provide a generalization of the main results, giving the defender another dimension of possible defense strategies, for example, by choosing time-varying variance of the noise, Σ_e . Finally, relaxing the assumption that the defender needs to know the noise model *a priori* would enable these results to be more applicable and is thus of high importance for future work.

ACKNOWLEDGMENTS

This research was funded by the Swedish Foundation for Strategic Research through the CLAS project (grant RIT17-0046).

REFERENCES

- [1] X. Liu and A. M. Eskicioglu. Selective encryption of multimedia content in distribution networks: challenges and new directions. In *Conf. Communications, Internet, and Information Technology*, pages 527–533, 2003.
- [2] C. Dwork and A. Smith. Differential Privacy for Statistics: What we Know and What we Want to Learn. *Journal of Privacy and Confidentiality*, 1(2), Apr. 2010.
- [3] F. Farokhi and H. Sandberg. Fisher Information as a Measure of Privacy: Preserving Privacy of Households With Smart Meters Using Batteries. *IEEE Transactions on Smart Grid*, 9(5):4726–4734, Sep. 2018.
- [4] C. Dwork. Differential Privacy: A Survey of Results. In *Theory and Applications of Models of Computation*, pages 1–19, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
- [5] J. Le Ny and G. J. Pappas. Differentially Private Filtering. *IEEE Transactions on Automatic Control*, 59(2):341–354, Feb 2014.
- [6] Z. Li and T. J. Oechtering. Privacy on hypothesis testing in smart grids. In *2015 IEEE Information Theory Workshop - Fall (ITW)*, pages 337–341, Oct 2015.
- [7] E. L. Lehmann and G. Casella. *Theory of Point Estimation*. Springer-Verlag, New York, NY, USA, second edition, 1998.
- [8] M. Basseville and I. Nikiforov. *Detection of Abrupt Change Theory and Application*, volume 15. 04 1993.
- [9] R. Alisic, M. Molinari, P. E. Par, and H. Sandberg. Bounding Privacy Leakage in Smart Buildings, 2020. <https://arxiv.org/abs/2003.13187>.
- [10] M. Jansen and G. Claeskens. Cramér Rao Inequality. In *International Encyclopedia of Statistical Science*, pages 322–323. Springer Berlin Heidelberg, 2011.
- [11] D. G. Chapman and H. Robbins. Minimum variance estimation without regularity assumptions. *Ann. Math. Statist.*, 22(4):581–586, 12 1951.
- [12] KTH Live-In Lab. <https://www.liveinlab.kth.se/>. Accessed: 2020-01-23.

- [13] IDA ICE 4.8 Software Program for Building Simulations. <https://www.equa.se/en/ida-ice>. Accessed: 2020-01-18.
- [14] J. B. Rawlings. Moving Horizon Estimation. In *Encyclopedia of Systems and Control*, pages 1–7. Springer London, London, 2013.
- [15] F. Farokhi and H. Sandberg. Ensuring privacy with constrained additive noise by minimizing Fisher information. *Automatica*, 99:275 – 288, 2019.

APPENDIX

Proof of Theorem 1. The estimator, $\psi_u(Y, M)$, uses the measurements Y and the model M in order to estimate the change time k^* . Here, k^* is a parameter which determines the probability distribution of Y . The minimum variance of the estimator is given by the Chapman-Robbins bound [11],

$$\text{Var}(\psi(Y, M)|k^*) \geq \sup_{\tau \neq 0} \frac{(\mathbb{E}[\psi(Y, M)|k^* + \tau] - \mathbb{E}[\psi(Y, M)|k^*])^2}{\mathbb{E}\left[\left(\frac{P(Y|k^* + \tau)}{P(Y|k^*)} - 1\right)^2 \middle| k^*\right]},$$

where $\tau \in \mathbb{Z}$ and $P(Y|k^*)$ is the probability of obtaining measurements Y , conditioned on the change time is k^* . Evaluating the expectation in the denominator gives

$$\text{Var}(\psi(Y, M)|k^*) \geq \sup_{\tau \neq 0} \frac{(\mathbb{E}[\psi(Y, M)|k^* + \tau] - \mathbb{E}[\psi(Y, M)|k^*])^2}{\int_{\mathbb{R}^N} \frac{P(Y|k^* + \tau)^2}{P(Y|k^*)} dY - 1}. \quad (15)$$

Since $\frac{P(Y|k^* + \tau)^2}{P(Y|k^*)} = e^{2\log P(Y|k^* + \tau) - \log P(Y|k^*)}$, we write for $\tau \geq 1$,

$$\begin{aligned} & \log P(Y|k^* + \tau) - \log P(Y|k^*) = \\ & -\frac{1}{2} \sum_{k=1}^N \left(y_k - CA^k x_0 - \sum_{l=k^* + \tau}^{k-1} CA^{k-l-1} Bu \right)^\top \Sigma_e^{-1} \left(y_k - CA^k x_0 - \sum_{l=k^* + \tau}^{k-1} CA^{k-l-1} Bu \right) + \\ & \frac{1}{2} \sum_{k=1}^N \left(y_k - CA^k x_0 - \sum_{l=k^*}^{k-1} CA^{k-l-1} Bu \right)^\top \Sigma_e^{-1} \left(y_k - CA^k x_0 - \sum_{l=k^*}^{k-1} CA^{k-l-1} Bu \right) = \\ & -\frac{1}{2} \sum_{k=k^* + 1}^N \left(\sum_{l=k^*}^{\min(k^* + \tau - 1, k-1)} CA^{k-1-l} Bu \right)^\top \Sigma_e^{-1} (2y_k - (2CA^k x_0 + \sum_{l=k^*}^{k-1} CA^{k-l-1} Bu + \sum_{l=k^* + \tau}^{k-1} CA^{k-l-1} Bu)). \end{aligned}$$

Continuing, we see that,

$$\begin{aligned} & 2\log P(Y|k^* + 1) - \log P(Y|k^*) = 2(\log P(Y|k^* + 1) - \log P(Y|k^*)) + \log P(Y|k^*) = \\ & -\frac{1}{2} \sum_{k=k^* + 1}^N \left(\sum_{l=k^*}^{\min(k^* + \tau - 1, k-1)} CA^{k-1-l} Bu \right)^\top \Sigma_e^{-1} (2y_k - (2CA^k x_0 + \sum_{l=k^*}^{k-1} CA^{k-l-1} Bu + \sum_{l=k^* + \tau}^{k-1} CA^{k-l-1} Bu)) - \\ & -\frac{1}{2} \sum_{k=1}^N \left(y_k - CA^k x_0 - \sum_{l=k^*}^{k-1} CA^{k-l-1} Bu \right)^\top \Sigma_e^{-1} \left(y_k - CA^k x_0 - \sum_{l=k^*}^{k-1} CA^{k-l-1} Bu \right) = \\ & -\frac{1}{2} \sum_{k=k^* + 1}^N (y_k - G)^\top \Sigma_e^{-1} (y_k - G) + \sum_{k=k^* + 1}^N \left(\sum_{l=k^*}^{\min(k^* + \tau - 1, k-1)} CA^{k-1-l} Bu \right)^\top \Sigma_e^{-1} \left(\sum_{l=k^*}^{\min(k^* + \tau - 1, k-1)} CA^{k-1-l} Bu \right) \\ & -\frac{1}{2} \sum_{k=1}^{k^*} \left(y_k - CA^k x_0 - \sum_{l=k^*}^{k-1} CA^{k-l-1} Bu \right)^\top \Sigma_e^{-1} \left(y_k - CA^k x_0 - \sum_{l=k^*}^{k-1} CA^{k-l-1} Bu \right), \end{aligned}$$

where,

$$G = CA^k x_0 + \sum_{l=k^*}^{k-1} CA^{k-l-1} Bu - 2 \left(\sum_{l=k^*}^{\min(k^* + \tau - 1, k-1)} CA^{k-1-l} Bu \right).$$

Inserting this expression into the bound in (15), evaluating the integral, and setting $\mathbb{E}[\psi(Y, M)|k^*] = k^* + g(k^*)$, we obtain

$$\text{Var}(\psi(Y, M)|k^*) \geq \max_{\tau \geq 1} \frac{(\tau + g(k^* + \tau) - g(k^*))^2}{e^{u^\top \mathcal{S}(\tau, M) u} - 1}, \quad (16)$$

where,

$$\mathcal{S}(\tau, M) = \sum_{k=k^* + 1}^N \left(\sum_{l=k^*}^{\min(k^* + \tau - 1, k-1)} CA^{k-1-l} Bu \right)^\top \Sigma_e^{-1} \left(\sum_{l=k^*}^{\min(k^* + \tau - 1, k-1)} CA^{k-1-l} Bu \right).$$

For $\tau \leq -1$, an equivalent calculation can be made giving the same expression as (16), but replacing \mathcal{S} with

$$\mathcal{S}_-(\tau, M) = \sum_{k=\hat{k}+\tau+1}^N \left(\sum_{l=\hat{k}+\tau}^{\min(\hat{k}-1, k-1)} CA^{k-1-l}Bu \right)^\top \Sigma_e^{-1} \left(\sum_{l=\hat{k}+\tau}^{\min(\hat{k}-1, k-1)} CA^{k-1-l}Bu \right).$$

However, since

$$\mathcal{S}(|\tau|, M) \leq \mathcal{S}_-(\tau, M),$$

for each positive integer $|\tau|$, we can ignore the $\tau \leq -1$ cases. □